

Amendments to the Claims:

This listing of claims will replace all prior versions, and listings, of claims in the application:

Listing of Claims:

1. (currently amended) A method of scheduling a media object for transmission between a server and a client, the method comprising:

partitioning the media object into segments, wherein the segments contain one or more blocks, wherein each block is a unit of media for which the client will wait to receive an entire block before playing out the block, and wherein each segment includes an integer number of blocks and wherein blocks are encoded into output symbols such that the output symbols generated are independent of when a client begins reception, wherein the output symbols are encoded such that ~~when the output symbols are transmitted to the client, the output symbols are transmitted they are transmittable to the client~~ in a manner that is independent of ~~the~~ which other output symbols were transmitted to the client;

determining one or more channels on which to serve each segment, the channels capable of carrying data between the server and the client;

determining a rate at which to serve each segment; and

determining a schedule pair for each channel, the schedule pair including a time at which the client may start receiving on the channel and a time at which the client may stop receiving on the channel.

2. (original) The method of claim 1 wherein if the client minimally fulfills the schedule pair for each channel, the client will be able to play out the media object uninterrupted after a startup latency.

3. (original) The method of claim 1 wherein the partitioning step includes first partitioning the media object into segments, then partitioning each segment into an integer number of blocks.

4. (original) The method of claim 1 wherein the partitioning step includes first partitioning the media object into blocks, then grouping the blocks into segments.

5. (original) The method of claim 1 wherein a maximum download rate at the client is less than an aggregate server transmission rate.

6. (previously presented) The method of claim 1 wherein the client is configured with a maximum download rate that does not limit the rate that each segment is served.

7. (original) The method of claim 1 wherein a maximum download rate at the client is only slightly greater than a media object play out rate.

8. (original) The method of claim 1 wherein a maximum download rate at the client is less than a media object play out rate.

9. (original) The method of claim 1 wherein at least one segment includes more than one block.

10. (original) The method of claim 1 wherein each segment includes exactly one block.

11. (original) The method of claim 1 wherein a plurality of blocks in a segment are a same size.

12. (original) The method of claim 1 wherein a plurality of blocks in a plurality of segments are a same size.

13. (original) The method of claim 1 wherein at least two of the channels are served at different rates.

14. (original) The method of claim 1 wherein at least two segments have different sizes.

15. (original) The method of claim 1 wherein each block in each segment is served at a same rate.

16. (original) The method of claim 1 wherein at least one block in at least one segment is served at a rate different than other blocks in the segment.

17. (original) The method of claim 1 wherein the segments are served at a rate less than a play out rate.

18. (original) The method of claim 1 wherein at least one segment is served on at least two channels, and wherein each of the at least two channels is served at a rate.

19. (original) The method of claim 18 wherein the each of the at least two channels are served at a same rate.

20. (original) The method of claim 18 wherein at least two of the each of the at least two channels are served at different rates.

21. (original) The method of claim 1 wherein the number of channels is less than or equal to a maximum number of concurrent channels at the client.

22. (original) The method of claim 2 wherein the client can minimally fulfill the schedule pair for each channel by downloading from a maximum number of concurrent channels.

23. (original) The method of claim 1 wherein the number of channels is less than or equal to a maximum number of concurrent channels served by the server.

24. (original) The method of claim 1 wherein the partitioning, determining one or more channels, determining a rate, and determining a schedule pair steps are performed so as to optimize a server bandwidth required to serve the media object.

25. (original) The method of claim 1 wherein the partitioning, determining one or more channels, determining a rate, and determining a schedule pair steps are performed so that the media object is served to the client at a rate less than or equal to a maximum client download rate.

26. (original) The method of claim 1 further comprising, for each of the plurality of segments, determining a size of the segment so that the segment is completely downloaded by the client prior to when the segment is due to be played out.

27. (original) The method of claim 26 wherein, for each of the plurality of segments, the rate at which to serve the segment is an integer multiple of a base rate.

28. (original) The method of claim 26 wherein all of the segments are served at a same rate.

29. (original) The method of claim 1 further comprising, for each of a plurality of segments:

if the segment may be scheduled to be served to the client without exceeding a maximum client download rate, scheduling the segment to be served to the client; and

if the segment may not be scheduled to be served to the client without exceeding the maximum client download rate, waiting to schedule the segment to be served until one or more other segments are scheduled to end being downloaded at the client.

30. (original) The method of claim 1 wherein the partitioning, determining one or more channels, determining a rate, and determining a schedule pair steps are performed so as to optimize a number of concurrent channels at the server required to serve the media object.

31. (original) The method of claim 1 wherein sizes of a plurality of blocks are each less than or equal to a maximum block size.

32. (original) The method of claim 1 wherein sizes of a plurality of segments are each less than or equal to a maximum segment size.

33. (original) The method of claim 32 wherein the maximum segment size is based on a maximum available storage at the client.

34. (original) The method of claim 1 wherein sizes of a plurality of blocks are each greater than or equal to a minimum block size.

35. (original) The method of claim 34 wherein the minimum block size is determined by an encoding scheme.

36. (original) The method of claim 34 wherein the minimum block size is determined by a media object player data format.

37. (original) The method of claim 1 wherein the partitioning step includes, for each of a plurality of segments, determining a start point of the segment and an end point of the segment based upon requirements of a media object player data format.

38. (original) The method of claim 1 wherein the partitioning step includes, for each of a plurality of blocks, determining a start point of the block and an end point of the block based upon requirements of a media object player data format.

39. (original) The method of claim 1 wherein the rate and the schedule pair are based on a maximum available storage at the client.

40. (original) The method of claim 1 wherein the rate varies over time.

41. (original) The method of claim 1 wherein at least one segment is served on at least two channels, and wherein a rate at which the segment is served on one of the at least two channels varies over time.

42. (original) The method of claim 1 wherein a set of channels serving at least one segment varies over time.

43. (currently amended) A system for scheduling a media object for transmission between a server and a client, comprising:

a module for partitioning the media object into segments, wherein the segments contain one or more blocks, wherein each block is a unit of media for which the client will wait to receive an entire block before playing out the block, and wherein each segment includes an integer number of blocks and wherein blocks are encoded into output symbols such that the output symbols generated are independent of when a client begins reception wherein the output symbols are encoded such that ~~when the output symbols are transmitted to the client, the output symbols are transmitted~~ they are transmittable to the client in a manner that is independent of the which other output symbols were transmitted to the client;

a module for determining one or more channels on which to serve each segment, the channels capable of carrying data between the server and the client;

a module for determining a rate at which to serve each segment; and

a module for determining a schedule pair for each channel, the schedule pair including a time at which the client may start receiving on the channel and a time at which the client may stop receiving on the channel.

44. (previously presented) A method of serving a media object, the method comprising:

receiving segments of a media object, wherein each segment includes an integer number of blocks, wherein each block is a unit of media for which a client will wait to receive an entire block before playing out the block, and wherein each block includes one or more input symbols;

for each segment, receiving an indication of one or more channels on which to serve the segment;

for each segment, receiving a rate at which to serve the segment;

determining an order in which to encode blocks;

generating output symbols for each block using a chain reaction code to generate the output symbols; and

transmitting the output symbols on the corresponding one or more channels, wherein each segment is served at the corresponding rate.

45. (original) The method of claim 44 further comprising providing one or more storage devices on which to store the media object.

46. (original) The method of claim 44 wherein the receiving a rate step includes receiving an aggregate rate at which to serve the segment on the one or more channels.

47. (original) The method of claim 44 wherein the receiving a rate step includes, for each of the one or more channels on which to serve the segment, receiving a rate at which to serve the segment on the each of the one or more channels.

48. (original) The method of claim 44 wherein a server serves at least two media objects concurrently.

49. (original) The method of claim 44 wherein a media object is concurrently served on at least two servers.

50. (original) The method of claim 49 wherein at least one segment is concurrently served on at least two servers.

51. (original) The method of claim 49 wherein at least one segment is served on only one server.

52. (original) The method of claim 44 wherein the media object is served on unicast channels.

53. (original) The method of claim 44 wherein the media object is served on multicast channels.

54. (original) The method of claim 44 wherein the media object is served on broadcast channels.

55. (canceled).

56. (original) The method of claim 44 wherein generating output symbols includes using a FEC code to generate output symbols.

57. (original) The method of claim 56 wherein all of the output symbols in a block are generated before transmitting any of the output symbols in the block.

58. (original) The method of claim 56 wherein for some or all of the plurality of output symbols, each of the some or all of the output symbols is generated upon a first transmission of the output symbol on the channel.

59. (original) The method of claim 56, further including determining an order in which to transmit output symbols corresponding to a block, and wherein transmitting output symbols includes transmitting output symbols corresponding to a block in the order.

60. (original) The method of claim 59 wherein the order is determined according to a random or pseudo-random sequence.

61. (original) The method of claim 44, further including determining an order in which to transmit output symbols for each block in a segment, and wherein transmitting output symbols includes, for each segment, transmitting output symbols corresponding to each block in the order.

62. (original) The method of claim 61 wherein the order is determined according to a random or pseudo-random sequence.

63. (original) The method of claim 44 wherein input symbols are used as output symbols.

64. (original) The method of claim 44 wherein a server transitions between serving a first media object and a second media object by successively stopping the serving of segments for the first media object and successively starting the serving of segments for the second media object.

65. (original) The method of claim 44 wherein the rate at least one segment is served on a channel varies over time.

66. (original) The method of claim 44 wherein at least one segment is served on at least two channels.

67. (original) The method of claim 44 wherein, for at least one segment, the one or more channels on which the segment is served varies over time.

68. (currently amended) An apparatus for serving a media object, comprising:

a block encoder coupled to receive segments of a media object, wherein each segment includes an integer number of blocks, wherein each block is a unit of media for which a client will wait to receive an entire block before playing out the block, wherein each block includes one or more input symbols; the block encoder including an input to receive an order in which to encode the blocks; and wherein the block encoder is configured to generate, in the order, output symbols for each block wherein blocks are encoded into output symbols such that the output symbols generated are independent of when a client begins reception, wherein the output symbols are encoded such that ~~when the output symbols are transmitted to the client, the output symbols are transmitted~~ they are transmittable to the client in a manner that is independent of ~~the~~ which other output symbols were transmitted to the client; and

a transmitter coupled to receive the output symbols from the block encoder, and coupled to receive, for each segment, an indication of one or more channels on which to serve the segment and a rate at which to serve the segment;

said transmitter configured to serve the output symbols on the corresponding one or more channels at the corresponding rate.

69. (currently amended) A method of receiving a media object that includes segments, wherein each segment includes an integer number of blocks, and wherein each block is a unit of media for which a client will wait to receive an entire block before playing out the block, the method comprising:

receiving a media object description of the media object;

joining and leaving each of a plurality of channels according to the media object description to download the blocks as a plurality of output symbols wherein the output symbols are encoded such that the output symbols need not depend on when joining and leaving occurs, wherein the output symbols are encoded such that ~~when the output symbols are transmitted to the client, the output symbols are transmitted~~ they are transmittable to the client in a manner that is independent of ~~the~~ which other output symbols were transmitted to the client;

reassembling the blocks in each segment; and

playing the blocks out in an order after a startup latency.

70. (original) The method of claim 69 wherein a client joins channels according to an order in the media object description.

71. (original) The method of claim 69 wherein, reassembling the blocks of a first segment and playing out the blocks of a second segment occur concurrently.

72. (original) The method of claim 69 wherein segments are downloaded on unicast channels.

73. (original) The method of claim 69 wherein segments are downloaded on multicast channels.

74. (original) The method of claim 69 wherein segments are downloaded on broadcast channels.

75. (original) The method of claim 69 wherein a plurality of segments are downloaded concurrently at an aggregate rate, and wherein the aggregate rate is less than a maximum download rate.

76. (original) The method of claim 69 wherein the media object is downloaded by a client at a rate less than an aggregate server transmission rate.

77. (original) The method of claim 69 wherein the media object is downloaded by a client at an unconstrained rate.

78. (original) The method of claim 69 wherein the media object is downloaded by a client at a rate only slightly greater than a media object play out rate.

79. (original) The method of claim 69 wherein the media object is downloaded by a client at a rate less than a media object play out rate.

80. (original) The method of claim 69 wherein joining and leaving a plurality of channels includes downloading at least a first segment upon the ending of the downloading of a second segment.

81. (original) The method of claim 69 wherein a rate at which the media object is downloaded may be adjusted by joining and leaving channels.

82. (original) The method of claim 81 wherein a client increases its reception rate by adding channels according to an order and decreases its reception rate by dropping channels either in a reverse of the order, or when a segment completes downloading on that channel.

83. (original) The method of claim 81 wherein a client increases its reception rate when it experiences no congestion and decreases its reception rate when it experiences congestion.

84. (original) The method of claim 69 wherein the blocks include original data, and wherein reassembling the blocks in each segment includes reordering the original data according to its temporal position in each block.

85. (original) The method of claim 69 wherein reassembling the blocks in each segment includes decoding with a FEC decoder.

86. (original) The method of claim 69 wherein reassembling the blocks in each segment includes decoding with a chain reaction decoder.

87. (original) The method of claim 69 wherein a single media object is downloaded from at least two servers.

88. (original) The method of claim 87 wherein at least one segment is downloaded from at least two servers.

89. (original) The method of claim 87 wherein at least one segment is downloaded from only one server.

90. (original) The method of claim 69 wherein a same segment in the media object is downloaded from at least two servers.

91. (original) The method of claim 69 further comprising playing out a pre-downloaded segment during the startup latency.

92. (original) The method of claim 69 wherein the media object is downloaded by a client from a maximum number of channels concurrently.

93. (original) The method of claim 69 wherein a set of channels on which at least one segment is downloaded varies over time.

94. (previously presented) A system for receiving a media object that includes segments, wherein each segment includes an integer number of blocks, and wherein each block is a unit of media for which a client will wait to receive an entire block before playing out the block, comprising:

a module for handling input of a media object description of the media object;

a module for handling channel joins and channel leaves for each of a plurality of channels according to the media object description, wherein the channels are capable of use for downloading the blocks to the client as a plurality of output symbols wherein the output symbols are encoded such that the output symbols need not depend on when joining and leaving occurs, wherein the output symbols are encoded such that ~~when transmitted to the client, the output symbols are transmitted~~ they are transmittable to the client in a manner that is independent of the ~~which other~~ output symbols ~~were~~ transmitted to the client;

a module for reassembling the blocks in each segment; and

a module for playing the blocks out in an order after a startup latency.

95. (previously presented) The method of claim 1, wherein the blocks are encoded into the output symbols according to a chain reaction code.